Augmentation of an Intelligent Flight Control System for a Simulated C-17 Aircraft

Karen Gundy-Burlet* and K. Krishnakumar.†
NASA-Ames Research Center, Moffett Field, CA 94035-1000

Greg Limes‡ and Don Bryant§
QSS Inc., Moffett Field, CA 94035-1000

This paper examines a neural-adaptive flight control system augmented with linear programming theory and adaptive critic techniques for a simulated C-17 aircraft. The baseline Intelligent Flight Control (IFC) system is composed of a neural network based direct adaptive control approach for applying alternate sources of control authority in the presence of damage or failures in order to achieve desired flight control performance. Neural networks are used to provide consistent handling qualities across flight conditions, adapt to changes in aircraft dynamics and make the controller easy to apply when implemented on different aircraft. In this study, IFC has been augmented with linear programming (LP) theory and adaptive critic technologies. LP is used to optimally allocate requested control deflections and the adaptive critic modifies the parameters of the aircraft reference model for consistent handling qualities. Full-motion piloted simulation studies were performed on a Boeing C-17. Subjects included NASA and Air Force pilots. Results, including subjective pilot ratings and time response characteristics of the system, demonstrate the potential for improving handing qualities and significantly increased survivability rates under various simulated failure conditions.

Nomenclature

B	=	control derivative matrix
e	=	error
f	=	generic function
J	=	cost function
k	=	reference model gain
M	=	pitching moment
p	=	roll rate
q	=	pitch rate
r	=	yaw rate
X	=	state vector
t	=	time
U	=	utility function
u	=	control vector
w	=	weighting vector
δ_a	=	aileron deflection
δ_e	=	elevator deflection
γ	=	discount factor

* Research Scientist, M/S 269-3, NASA Ames Research Center, Moffett Field, CA, 94035. Associate Fellow, AIAA.
† Research Scientist, M/S 269-1, NASA Ames Research Center, Moffett Field, CA, 94035. Associate Fellow, AIAA.
‡ Computer Engineer, M/S 269-1, QSS. Inc., Moffett Field, CA, 94035.
§ Pilot, M/S 269-3, QSS. Inc., Moffett Field, CA, 94035

λ = control reallocation matrix
ω = reference model frequency

I. Introduction

In the last 30 years, aircraft flight control system failures have claimed more than a thousand lives. These accidents were typically related to jammed cables, faulty valves or structural failures leading to lost hydraulic systems. In particular for the United Airlines 232 accident[1], an uncontained engine failure severed all three redundant hydraulic lines leaving the pilots with only manually operated throttles to control the aircraft. The aircraft crash-landed, but the pilots' efforts to control the plane using alternate sources of control authority saved many lives. As a result, Burcham, et al.[2] developed control algorithms to significantly reduce pilot effort in utilizing alternate sources of control authority, and successfully flew a F-15[2] and eventually an MD-11[3] under propulsion-only control.

Rysdyk and Calise[4] and Kim and Calise[5] developed a fault-tolerant neural flight control architecture later utilized by Kaneshige and Gundy-Burlet[6] for demonstration of the Integrated Neural Flight and Propulsion Control System (INFPCS) on transport class aircraft. The neural network based approach incorporates direct adaptive control with dynamic inversion to provide consistent handling qualities without requiring extensive gain-scheduling or explicit system identification. It also uses an a-priori hierarchical control allocation technique to ensure that conventional flight control surfaces are utilized under normal operating conditions. Under damage or failure conditions, the system re-allocates control to non-traditional flight control surfaces and/or incorporates propulsion control, when additional control power is necessary for achieving desired flight control performance.

The research reported in this paper is an extension of the above work to include a full simplex method LP technique for control reallocation[7] and adaptive critic technologies[8] for reference model adaptation. The LP technique is utilized to optimally reallocate control across a potentially diverse list of surfaces, where it is difficult to manually form an a-priori hierarchical list. In cases where the system is degraded to a level where control authority is insufficient to achieve the required reference model characteristics, an adaptive critic is used to degrade the reference model to the capabilities of the airplane. These algorithmic techniques are considered to be at a higher level of intelligence in the hierarchy of intelligent control[9,10]. A sample of other adaptive critic analysis and applications can be found in references 12-14.

This paper presents a brief overview of levels of intelligent control and discusses the elements of the augmented IFC system architecture in that context. It then provides overviews of the linear programming control allocation technique and adaptive critic algorithms. Piloted simulation studies were performed at NASA Ames Research Center on a full motion simulator utilizing a model of a Boeing C-17. Subjects included NASA and Air Force test pilots. Results include handling quality comparisons, landing performance and time histories comparing the current augmented IFC performance to that of the baseline IFC system under nominal and simulated failure conditions.

II. Levels of Intelligent Control

Over the past decade, several innovative control architectures utilizing the intelligent control tools have been proposed. KrishnaKumar9,10, has proposed a classification scheme based on the ability of the intelligent flight control architecture for self-improvement (see Table 1). The classification scheme divides the control architectures among levels of intelligent control. For instance, most of the proposed architectures can be divided among level 0, level 1, level 2, and level 3 intelligent control schemes. Based on this classification scheme, several seemingly differing control architectures can be seen as achieving similar goals.

. **Table 1. The Levels of Intelligent Control**

Level	Self improvement of	Description
0	Tracking Error (TE)	Robust Feedback Control: Error tends to zero.
1	TE + Control Parameters (CP)	Adaptive Control: Robust feedback control with adaptive control parameters (error tends to zero for non-nominal operations; feedback control is self improving).
2	TE + CP + Performance Measure (PM)	Optimal Control: Robust, adaptive feedback control that minimizes or maximizes a utility function over time.
3	TE+CP+PM+ Planning Function	Planning Control: Level 2 + the ability to plan ahead of time for uncertain situations, simulate, and model uncertainties.

III. High Level System Architecture

Fig. 1 presents a block diagram of the Level 2 IFC on the C-17 test bed used in this study. The levels of intelligent control outlined earlier are labeled in the figure. It should be noted that Level 0 is non-adaptive whereas Levels 1 and 2 are adaptive. Level 1 is non-optimal whereas Level 2 is optimal. Details of each block in the figure are presented below.

Editor…please insert a hyperlink to Fig.1.pdf here.

Fig. 1 Neural-Adaptive Flight Control System Architecture.

A. Reference Models (Level 0)

The pilot commands roll and pitch rates and aerodynamic lateral accelerations through stick and rudder pedal inputs. These commands are then transformed into body-axis rate commands, which also include turn coordination, level turn compensation, and yaw-dampening terms. First-order reference models are used to filter these commands in order to shape desired handing qualities.

B. Dynamic Inversion/Aero Generation (Level 0)

The dynamic inversion element[15] converts the summed response commands into virtual control surface commands. Dynamic inversion is based upon feedback linearization theory. No gain-scheduling is required, since gains are functions of aerodynamic stability and control derivative estimates and sensor feedback. Several methods are available to accomplish approximate model definition: simple linear model methods, nonlinear tables or using pre-trained neural networks (non-changing) to provide estimates of aerodynamic stability and control characteristics. The model is then inverted to solve for the necessary control surface commands. In our work, a Levenberg-Marquardt (LM) multi-layer perceptron[16] is used to provide dynamic estimates for model inversion. The LM network is pre-trained with stability and control derivative data generated by a Rapid Aircraft Modeler, and vortex-lattice code[17].

C. P + I Error Controller (Level 0)

Errors in roll rate, pitch rate, and yaw rate responses can be caused by inaccuracies in aerodynamic estimates and model inversion. Unidentified damage or failures can also introduce additional errors. In order to achieve a rate-command-attitude-hold (RCAH) system, a proportional-integral (PI) error controller is used to correct for errors detected from roll rate, pitch rate, and yaw rate (p, q, r) feedback. The root cause of the errors are not specifically identified and may come from many factors including damage to the aircraft, poor estimates of stability and control derivatives or through the assumed linearized model of the system.

D. Learning Neural Network (Level 1)

The on-line learning neural networks work in conjunction with the error controller. By recognizing patterns in the behavior of the error, the neural networks can learn to remove biases through control augmentation commands. These commands prevent the integrators from having to windup to remove error biases. By allowing integrators to operate at nominal levels, the neural networks enable the controller to provide consistent handling qualities. The learning neural networks not only help control the nominal system, but also provide an additional potential for adapting to changes in aircraft dynamics due to control surface failures or airframe damage. A Lyapunov stability proof guarantees boundedness of the tracking error and network weights.[4]

E. Optimal Allocation (Level 2)

This system uses a linear programming technique to optimally allocate required rate commands to available control surfaces based on perceived limits. Control derivatives for each axis were roughly estimated for every available control surface on the aircraft via the Integrated Vehicle Modeling Environment17. The LP solver utilized this information in conjunction with a cost function to determine the distribution of surface deflection commands. The cost function biased the solution toward the minimum drag configuration and the smallest possible surface deflections to achieve the desired rates. Structural limitations for the subject aircraft are not known, and were not incorporated into the cost function, but the technique admits their potential inclusion in the future. Unconventional flight control surface allocations are only utilized when the primary flight control surface commands exceed the known limits of deflection. For example, yaw rate control is normally provided through rudder deflection. If this command should saturate, then the remaining portion of the command is applied via a blended solution that could result in the deflection of ailerons and/or spoilers.

F. Adaptive Critic (Level 2)

In the event of a severe degradation in performance of an aircraft, pilot handling qualities as dictated by the reference model cannot be maintained. It will be desirable to "optimally" modify the dynamics of the reference model to suit the situation in hand. Towards these goals an adaptive critic is utilized to optimize the shape of the reference model dynamics in the event of a failure or damage.

IV. Control Allocation Detail

In the material that follows, we outline the three different control allocation techniques used in this study. These are: 1) Daisy chain control allocation: 2) Optimal control allocation using linear programming: 3) A table-look up with blending.

A. Daisy Chain

In the daisy chain approach[6], secondary control surfaces are used in a hierarchical form to compensate for failures in primary control surfaces. Propulsion control was not used in this study, but if utilized, the engines become tertiary control effectors. The table below presents the hierarchy. The alternate control sources are used only when the limits of the primary control surfaces are exceeded.

Table 2. Daisy Chain Allocation Table

	Elevator	Symmetric Aileron	Differential Aileron	Rudder
Pitch Axis	Primary	Secondary		
Roll Axis			Primary	Secondary
Yaw Axis				Primary

B. Optimal Control Allocation

In optimal control allocation, the choice of the hierarchy is not predetermined as in daisy chain. The choice of alternate surfaces depends on both surface effectiveness and a cost function. The cost function is used to bias the solution toward configurations of interest (i.e. minimum drag). In this section, we present the equations that lead to a problem formulation that is tractable to optimize in real time using linear programming.

The aircraft dynamic system is conveniently defined as:

$$\left[\dot{X}\right] = f(X) + \left[B\right]\left[u\right] + f_{trim} \tag{1}$$

where X is the state vector, B is the control derivative matrix and u is the control vector. Let a portion of the vector u hit the limit u_L. We now partition the $\left[B\right]\left[u\right]$ matrix as follows:

$$\begin{bmatrix} B_{UU} & B_{UL} \\ B_{LU} & B_{LL} \end{bmatrix}\begin{bmatrix} u_U \\ u_L + \Delta u_L \end{bmatrix} = \begin{bmatrix} B_{UU} & B_{UL} \\ B_{LU} & B_{LL} \end{bmatrix}\begin{bmatrix} u_U \\ u_L \end{bmatrix} + \begin{bmatrix} B_{UU} & B_{UL} \\ B_{LU} & B_{LL} \end{bmatrix}\begin{bmatrix} 0 \\ \Delta u_L \end{bmatrix} \tag{2}$$

where u_U = vector of UNLIMITED control variables and $u_L + \Delta u_L$ = vector of LIMITED control variables, with u_L =limit. To overcome the control needed beyond the limit, defined as Δu_L, we need the unlimited control variables u_U to compensate by an amount Δu_U. The needed relationship to compute Δu_U is:

$$\begin{bmatrix} B_{UU}\Delta u_U \\ B_{LU}\Delta u_U \end{bmatrix} = \begin{bmatrix} B_{UL}\Delta u_L \\ B_{LL}\Delta u_L \end{bmatrix} \tag{3}$$

Let us now define a control reallocation matrix $\left[\lambda\right]$ such that $\left[\Delta u_U\right] = \left[\lambda\right]\left[\Delta u_L\right]$. Substituting this into Eq. 3, we get $\begin{bmatrix} B_{UU} \\ B_{LU} \end{bmatrix}\left[\lambda\right] = \begin{bmatrix} B_{UL} \\ B_{LL} \end{bmatrix}$

For convenience, let us define the above relationship in a compact matrix-vector form: $\left[\alpha\right]\left[\lambda\right] = \left[\beta\right]$ or equivalently

$$\left[\alpha\right]\left[\lambda_1 \quad \lambda_2 \quad . \quad . \quad \lambda_m\right] = \left[\beta_1 \quad \beta_2 \quad . \quad . \quad \beta_m\right] \tag{4}$$

where

$\lambda_1 \quad \lambda_2 \quad . \quad . \quad \lambda_m$ are the columns of matrix $\left[\lambda\right]$

$\beta_1 \quad \beta_2 \quad . \quad . \quad \beta_m$ are the columns of matrix $\left[\beta\right]$

To allocate the available control surfaces optimally, we define the following Linear Programming problem:

$$\min_{w_i} \left(w_i^T \lambda_i\right) \tag{5}$$

subject to $\left[\alpha\right]\left[\lambda_i\right] = \left[\beta_i\right]$ and $\lambda_{\min} \le \lambda_{ij} \le \lambda_{\max}$

The weighting vector w_i can be used to effect the preferred choice of control reallocation. For example, to use ailerons first when elevators saturate, the corresponding elements of *w* for ailerons will be small and elements corresponding to other control surfaces will be big.

The optimization problem stated above might not be realizable in cases in which sufficient control authority does not exist after limit violation. A general framework for "standard-LP" based control allocation can now be stated as: $\min_{w_i} (w_i^T \lambda_i)$ subject to $[\overline{\alpha}][\lambda_i] \le [\overline{\beta}_i]$ and $\lambda_{ij} \ge 0$. In this equation, $[\overline{\beta}_i]$ = abs($[\beta_i]$) and $[\overline{\alpha}]$ represents the row elements of $[\alpha]$ have their signs reversed relative to the sign of the elements in $[\beta_i]$.

C. Table Driven Control Allocation with Blending

The look-up table concept was explored because of the perception that it would potentially be easier to certify for flight than an online LP algorithm. In addition, one of the problems with adapting to system faults with no explicit fault identification is the time needed for the adaptive element to wind-up the error to drive the failed control surface command to its perceived limit. This "dead band" translates to a degraded (highly non-linear) handling quality of the aircraft. To alleviate this, one could blend primary and secondary controls just before the perceived surface limit is reached. This can be achieved by defining certain pseudo limits for each of the surfaces to achieve early transition to secondary controls in anticipation of hitting the actual limits. In the next set of equations, we present the equations for implementing the blending for the pitch axis. Similar equations could be derived for roll and yaw axes.

Assuming that $\frac{\partial M}{\partial \delta_e}$ and $\frac{\partial M}{\partial \delta_a}$ are the pitching moment derivatives due to elevators and ailerons and are known and Pseudo limits $= \delta_{e1}$ are given (similarly for other axes), then the net effective pitch acceleration due to control, after pseudo limit has been hit is computed as:

$$k \frac{\partial M}{\partial \delta_e} \delta_e = \frac{\partial M}{\partial \delta_e} \delta_{e1} + \left[\alpha \frac{\partial M}{\partial \delta_e} (\delta_e - \delta_{e1}) \right]_{\text{Primary}} + \left[\beta \frac{\partial M}{\partial \delta_e} (\delta_e - \delta_{e1}) \right]_{Secondary} \tag{6}$$

To achieve blending, the tunable gains, β, α, that are limited to lie between 0 and 1, need to be set. If $\beta = 0,\ \alpha = 1$, we get no blending. If $\beta = (1 - \alpha),$ we get blending with unity gain ($k = 1$). If $\beta > (1 - \alpha),$ we get $k > 1$ and there is an effect of amplifying the controller gain.

For pitch control, we can compute the aileron needed to blend for the required pitch acceleration using Eq. 6 as follows:

$$\frac{\partial M}{\partial \delta_a} \Delta \delta_a = \left[\beta \frac{\partial M}{\partial \delta_e} (\delta_e - \delta_{e1}) \right] \tag{7}$$

Rearranging, we get

$$\Delta \delta_a = \beta (\delta_e - \delta_{e1}) \frac{\partial M / \partial \delta_e}{\partial M / \partial \delta_a} \tag{8}$$

A look-up table can be then constructed for each axis using similar derivations. For instance, the entries of the look-up table for blending aileron control for the pitch axis will be the constants β, δ_{e1}, and $\frac{\partial M / \partial \delta_e}{\partial M / \partial \delta_a}$.

V. Adaptive Critic Detail

Adaptive critic designs have been defined as designs that attempt to approximate dynamic programming based on the principle of optimality. Adaptive critic designs consist of two entities, an action network that produces optimal actions and an adaptive critic that estimates the performance of the action network. The adaptive critic is an optimal or near optimal neural network estimator of the cost-to-go function that is trained (adapted) using recursive equations derived from dynamic programming. The critic is termed adaptive as it adapts itself to output the optimal cost-to-go function from a given system state. The action network is adapted simultaneously based on the information provided by the critic. The action network consists of any piece of the overall control architecture that has an effect on the final performance of the closed-loop system. In typical applications, the action network consists of the controller that is optimized using the critic. The inputs required for designing an adaptive critic design are

- The cost function or the performance measure.
- A parameterized representation of the critic.
- A parameterized representation of the action network.
- A method for adapting the parameters of the critic.

D. Choice of the cost function

The choice of the cost function comes from the problem at hand. The cost could be distributed over the entire length of time or be defined at the end of the process. Typical examples of the two types are minimizing the fuel spent for a certain flight mission or intercepting a projectile where the utility depends only on the final error. Typically, the cost function can be given as,

$$J = \sum_{i=0}^{T} \gamma^i U[x(i), u(i)] \tag{9}$$

where $U[x(i), u(i)]$ is the utility function or a penalty function that is a function of the state of the system, *x(i)*, and the control (action), *u(i)*, given to the system. 'γ' is a discount factor that discounts the future performance.

The dynamic programming principle states that we can formulate an optimal control problem where we can get an optimal solution by minimizing the cost-to-go function, *J(t),* which is defined as,

$$J(t) = \sum_{i=1}^{T-t} \gamma^i U[x(t+i), u(t+i)] \tag{10}$$

So the critic is designed to approximate the optimal form of this cost-to-go function or its derivatives with respect to the state of the system depending on the particular adaptive critic design.

E. Parameterized representation for the critic and the action network

Parameterization of the critic and the action network is achieved by the use of neural networks. Having learnt to model a system, neural networks can be used to provide sensitivities of the system outputs with respect to the system inputs. This proves to be useful information especially for the training of our intelligent control architecture. Reference 6 provides detailed insight into the area of neural networks and their use in control.

F. Training the critic

Several methods have been proposed for training adaptive critics that are based on the dynamic programming equation. These methods vary based on the level of complexity and the degree of accuracy sought for training the cost-to-go function. Some of these methods are the heuristic dynamic programming (HDP) approach, the dual heuristic programming (DHP) approach, and the global dual heuristic dynamic (GDHP) programming approach. The critic is implemented as a Dual-heuristic Dynamic Programming (DHP) critic. The DHP scheme is similar in idea to

the HDP scheme, however, in DHP, the critic outputs the derivative of the performance with respect to control directly, which is the signal necessary for adapting the controller. Though DHP is more complex both theoretically and in implementation, it is generally considered to produce better results. References 11-12 provide a more detailed discussion on the subject.

G. Adaptive Critic Application

A static reference model is sufficient when the system is functioning normally. The model needs to change when desired performance is not achievable with the available control authority. If this problem is not rectified, two issues arise: (1) wrong signal for NN training for Level 1; (2) error integrator wind-up. The adaptive critic application to reference model adaptation addresses these issues. Another issue, although not considered in this paper, is the use of engines for rotational control. Engine responses are not fast enough to provide the same handling qualities. In cases where propulsion control is used, adaptive critics can be used to optimally adjust the reference model frequencies.

H. Use of the Reference Model as the Action Network

The architecture of the adaptive critic block from Fig. 1 is shown in detail for the aircraft control problem in Fig. 2. In many of the methods using adaptive critics, the trained critic is finally used to update the controller. In this implementation, the reference model is looked at as the action network. In other words, the reference model is the input into the closed-loop system consisting of the Level 1 IFC and thus is the right choice as the action network. In the face of any failures or damage, it is sometimes impossible for the controller to achieve the system outputs as demanded by the original reference model. The Level 2 IFC using the adaptive critic neural network therefore attempts to adapt the parameters of the reference model to provide realizable performance requirements on the system.

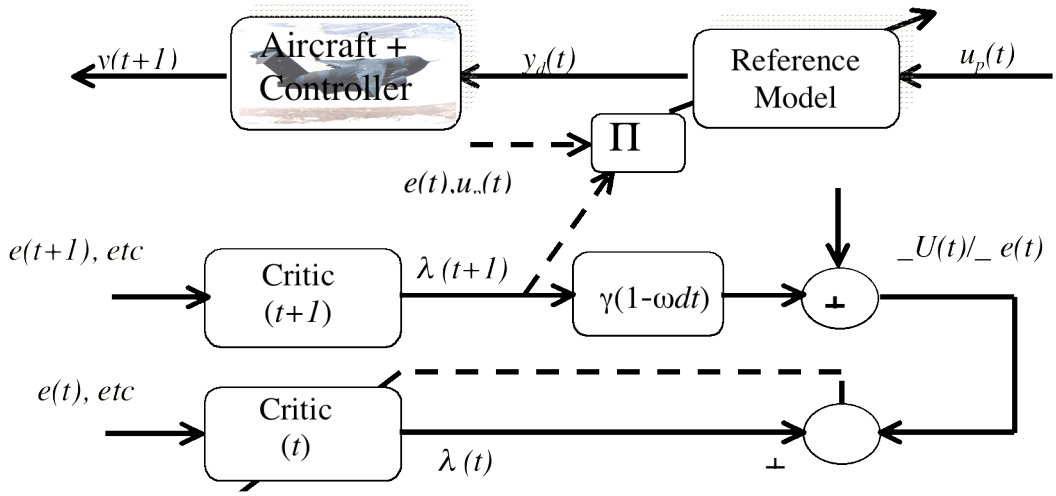

Fig. 2. **Reference model tuning using the adaptive critic**

I. Adaptive Critic Neural Network

Adaptation is achieved using a single hidden-layer neural network with the following inputs:

- error in the corresponding axes at time *t*
- error rate in the corresponding axes at time *t*
- control needed beyond the limits at time *t*
- pilot inputs at time *t*

The neural network output is defined to be the derivative of the performance index with respect to the individual error in the axes (roll rate error, pitch rate error, and yaw rate error). The pertinent equations are given below:

$$J(t) = \sum_{t=T}^{\infty} U(t) \quad \Rightarrow \quad \lambda_e(t) = \left(\frac{\partial J}{\partial e}\right)_t$$
$$U(t) = f(e^2)$$
$$f(e^2) = \frac{1}{1 + \exp(-m(e^2 - c))} \tag{11}$$

where e = error in each of the three axes (p, q, r) and m, c = constants chosen by the user. A sigmoid function is chosen as a penalty function since it penalizes the performance measure only if the error rises beyond a certain value, c. The constant m defines the slope of the Sigmoidal curve. It can be seen from Eq. 11, for a value of $c = 0.01$, the performance measure starts getting penalized when the error approaches 0.1 rad/sec. So, for values less than 0.1 rad/sec the penalty is close to zero and for values beyond 0.1 rad/sec the penalty is close to one. At the same time since the penalty function goes from zero to one, it automatically provides itself as a normalized function.

J. DHP Critic Training Equations

We use a first order reference model for our problem. Given

$$\dot{y}_d = -\omega_d y_d + k_d \omega_d u_p, \tag{12}$$

subtracting $\dot{y}$ and $\omega_d y$ from both sides of the equation and rearranging we get (Note: $y = actual\ \ p, q, r$).

$$\dot{y} - \dot{y}_d = -\omega_d (y - y_d) - k_d \omega_d u_p + (\dot{y} + \omega_d y)$$
$$\dot{e} = -\omega_d e - k_d \omega_d u_p + d\hat{\delta} \tag{13}$$

In the above equation, $d\hat{\delta} = (\dot{y} + \omega_d y)$ represents the additional control needed and could be seen in this context as an external disturbance that causes error to behave non-optimally.

In discrete form (first order forward differencing),

$$e_{t+1} - e_t = (-\omega_d e_t - k_d \omega_d u_{p,t} + d\hat{\delta})dt$$
$$e_{t+1} = (1 - \omega_d dt)e_t - (k_d \omega_d dt)u_{p,t} + \dot{y}dt + \omega_d y dt \tag{14}$$

From Dynamic Programming, we know that $\min_{\omega_d, k_d} J_t = \min_{\omega_d, k_d} (\gamma\ J_{t+1} + U_t)$ where γ is a discount factor (usually > 0.9) Differentiating with respect to e_t and noting that minimization is an approximation and that J is only an estimate, we have

$$\left(\frac{\partial \hat{J}_t}{\partial e_t}\right)_{desired} = \left[\hat{\lambda}_t\right]_{desired} = \gamma\ \frac{\partial \hat{J}_{t+1}}{\partial e_t} + \frac{\partial U_t}{\partial e_t} \tag{15}$$

with

$$\frac{\partial U_t}{\partial e_t} = \frac{2 * m * e * \exp(-m(e^2 - c))}{(1 + \exp(-m(e^2 - c)))^2} \tag{16}$$

and

$$\frac{\partial \hat{J}_{t+1}}{\partial e_t} = \frac{\partial \hat{J}_{t+1}}{\partial e_{t+1}} \frac{\partial e_{t+1}}{\partial e_t} = \left[\hat{\lambda}_{t+1}\right](1 - \omega_d dt) \tag{17}$$

Implying

$$\left(\frac{\partial \hat{J}_t}{\partial e_t}\right)_{desired} = \left[\hat{\lambda}_t\right]_{desired} \qquad = \gamma \left[\hat{\lambda}_{t+1}\right](1 - \omega_d dt) + \frac{\partial U_t}{\partial e_t} \tag{18}$$

The above quantity is the training signal for the critic at time "t"

K. Reference Model Adaptation Equations

Now we derive the computation of the performance sensitivities, ($\frac{\partial J}{\partial \omega_{d1}}$ and $\frac{\partial J}{\partial k_d}$), for adapting the reference model parameters, ω_{d1}, k_d. In the rest of the presentation, the notation '^' has been dropped for convenience. Then:

$$\begin{aligned} &\frac{\partial J_{t+1}}{\partial \omega_d} = \left[\frac{\partial J_{t+1}}{\partial e_{t+1}}\right]\left(\frac{\partial e_{t+1}}{\partial \omega_d}\right) \\ &= -\lambda_{t+1} f_t dt \\ &where \\ &f_t = (e_t - y_t + k_d u_{p,t}) \end{aligned} \tag{19}$$

One can write a dynamic equation for the quantities under summation as follows,

$$S_{t+1} = f_t dt, \quad \infty \le t \le 1 \tag{20}$$

hence,

$$\frac{\partial J_{t+1}}{\partial \omega_d} = -\lambda_{t+1} S_{t+1} \text{ and } \frac{\partial J_{t+1}}{\partial k_d} = -\lambda_{t+1} R_{t+1} \tag{21}$$

where $R_1 = \omega_d u_{p,0} dt$ and $R_{t+1} = (1 - \omega_d) R_t + \omega_d u_{p,t} dt$

A smoothing algorithm can be used to desensitize the adaptation to sensor noise and other unwanted variations in the error estimate. This is achieved using a smoothing algorithm given below:

$$\frac{\partial J}{\partial \omega_d} = -\frac{1}{n+1}\sum_{t=t_{now}}^{t_{now}-n*dt} \lambda_t e_{t-dt}\, dt \quad \text{and} \quad \frac{\partial J}{\partial k_d} = -\frac{1}{n+1}\sum_{t=t_{now}}^{t_{now}-n*dt} \lambda_t u_{p,t-dt}\, dt \tag{22}$$

Now the parameters of the reference model can be adapted using the following gradient descent equations

$$\omega_d = \omega_d - \eta_\omega \frac{\partial J}{\partial \omega_d} + b_1, \quad \omega_{dL} \le \omega_d \le \omega_{dU}$$
$$k_d = k_d - \eta_k \frac{\partial J}{\partial k_d} + b_2, \quad k_{dL} \le k_d \le k_{dU} \tag{23}$$

where,

- η_ω, η_k = adaptation constants.
- b_1, b_2 = learning biases.
- $\omega_{dU}, \omega_{dL}, k_{dU}, k_{dL}$ = Upper and lower limits of variation.

The upper limit is set close to the normal gain and frequency of the reference model while the minimum values were set to half of the upper values. Too much lower than half the upper value led to an aircraft with exceedingly sluggish response and it was felt that it was better to "overdrive" the aircraft response and risk some uncomfortable behavior than deal with an unresponsive aircraft. Learning biases are used to recover the original frequencies and gains when the failure is removed. Also, by limiting the frequency and the gain, stability of the reference model is maintained.

L. An Optimistic Retrospective Critic (ORC)

As a first approximation, it is assumed that the critic's adaptation has progressed in the right direction and that the performance will achieve a stationarity condition in the next time step of interest. This is equivalent to saying that $\lambda_{t+1} = 0$. Using this assumption, the λ's are computed for preceding n time steps using Eq. 15 and the reference model is adapted using the rest of the equations through Eq. 23. This approach is named here as the Optimistic Retrospective Critic (ORC) to reflect the fact the approach is both optimistic (assuming $\lambda_{t+1} = 0$) and looks at n time steps into the past to compute the corrected λ's. The ORC approach can be seen as the bench-mark by which the adaptive critic performance can be evaluated.

.

VI. Test Articles and Facilities

The C-17 airplane is a high performance military transport with a quad-redundant Fly-by-Wire Flight Controls System. The flexibility of its control architecture makes it a suitable platform for various types of research in support of safety initiatives. As shown in Fig. 3, the aircraft has a stabilizer, four elevators, two ailerons, eight spoiler panels and two rudders, a total of 17 surfaces used for active flight control by the IFC systems. The aircraft also has 8 mechanically interconnected slats and 4 flaps for a total of 22 control surfaces and 4 engines on the aircraft.

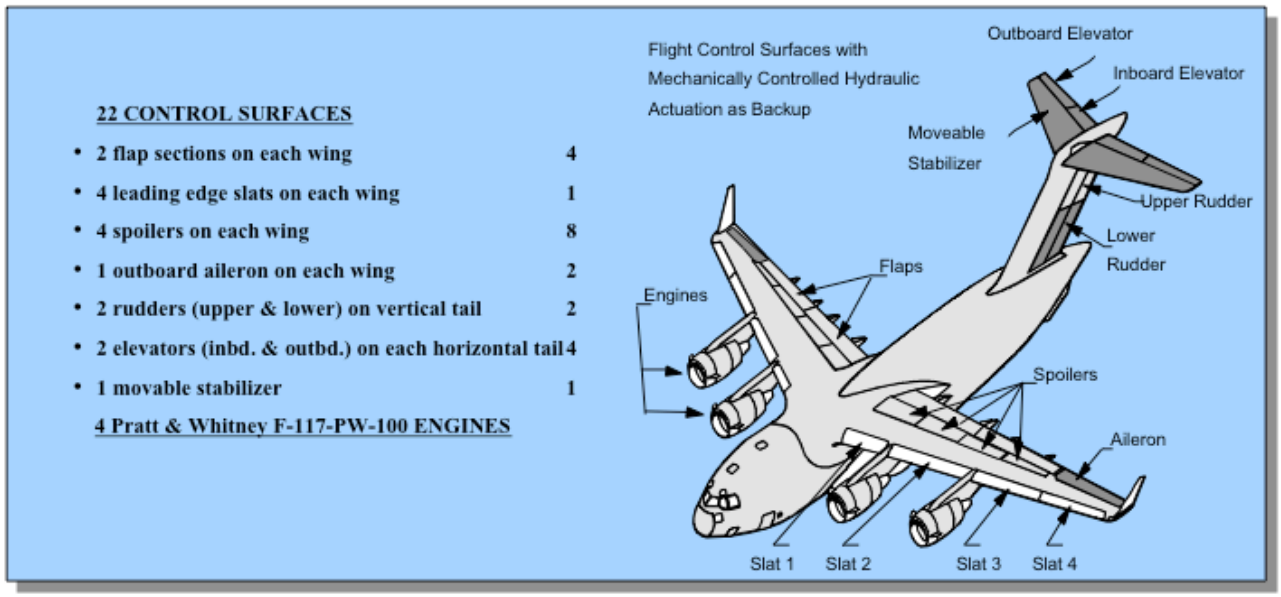

Fig. 3**. Boeing C-17 Control Surfaces.**

The level 2 IFC was evaluated utilizing the Crew Vehicle Systems Research Facility (CVSRF) at NASA-Ames Research Center. The Advanced Concepts Flight Simulator (ACFS)[18] has been modified to accommodate a model of a Boeing C-17 aircraft. The simulator is equipped with a six degree-of-freedom motion system, programmable flight displays, digital sound and aural cueing system, and a 180-degree field of view visual system.

M. Piloted results using the ORC

The goal of the C-17 experiment was to evaluate the performance of the Level 2 IFC system relative to that of the Level 1 IFC system. In addition to testing the system using time response characteristics, five C-17 pilots from the Air Force, and NASA were used to compare the IFC systems. For the piloted study, it was decided to use the ORC approach for its simplicity and to compare different control allocation strategies. The native C-17 flight control system was not included in the evaluation because an early version of the flight control system is incorporated into the simulation, and it was felt that it does not adequately represent the current C-17 SCAS. The evaluation criteria for all the pilots included Cooper-Harper[19] (CH) ratings, approach performance time history data, touchdown snapshot data and pilot comments.

The failure scenarios for the C-17 test are outlined in Table 3. The scenarios were designed to test performance of the controllers relative to primary failures in all 3 axes' as well as a failure sequence that would couple all the axes. The roll and pitch axis scenarios were utilized during normal landing operations, the yaw axis failure on a takeoff sequence, and the coupled failure during a tactical descent scenario. Asymmetric failure scenarios were randomly assigned to either the left or the right side of the plane. The pilots were asked to perform handling qualities tests consisting of roll, pitch and yaw doublets and provide individual CH ratings for each of the three axes.

Table 3**. Failure scenario characteristics.**

C17 Scenario	Scenario Characteristics: Winds 190 @ 10, light turbulence
Pitch axis	Full tail failure. Stabilizer failed at trim. 2 rudders, elevator*s* failed at 0 deg., in flight
Roll axis	2 ailerons and 8 spoiler panels failed at 0 deg., in flight
Yaw axis	Two engines out on one side on takeoff, minimum climb speed + 10Kts.
Coupled failure	During tactical descent (failures on one side) 23,000' : Stab frozen at trim 20,000' : 2 Elevators frozen at 0 deg. 17,000' : Upper rudder hard over 15,000' : Outboard flap fails retracted 14,000' : Aileron frozen at 0 deg. 13,000' : Two outboard spoilers frozen at 0 deg. When engines come out of reverse: Outboard engine seizes.

Fig. 4 shows delta Cooper Harper Ratings for Level 2 IFC systems (AKA Gen-3 in the legend) relative to Level 1 IFC system (AKA Gen-2). These were computed by subtracting each pilot's level 1 IFC CH ratings from the Level 2 systems ratings. Scenarios represented from left to right are: tail failure, wing failure, two-engine out takeoff and tactical descent scenario. Longitudinal and lateral CH ratings are reported separately due to the potentially extreme performance differences between the axes. In each block are shown the ratings for the pure LP solution and for the table driven control allocation scheme. The control allocation table was hand tuned to blend control across surfaces to minimize dead-bands and to eliminate extreme asymmetric deflections of the elevators for the purpose of minimizing stress on the tail. The order in which each control system was flown by the test pilot was randomized to minimize biases induced by pilot familiarity with the scenario, and that order is shown to the right of each rating delta. The data showed no clear trend with respect to test order. The average delta is listed over the top of each of the bars. As can be seen, the average delta showed improvement (lower CH rating) in most cases. The small degradation shown for the longitudinal tail failure case is probably related to the tuning of the Level 1 controller for longitudinal failure scenarios. Particular improvement is shown for the wing failure case, where the Level 1 controller only utilized yaw-based roll control as compared to the split elevator and rudder control utilized by the Level 2 controllers. The ratings for the two different Level 2 control allocation techniques depended on the importance of the control dead-bands to the handling qualities and suggests that either more control surface blending should be employed or that the controller should be integrated with a vehicle health monitor to remove failed surfaces from the list utilized by LP.

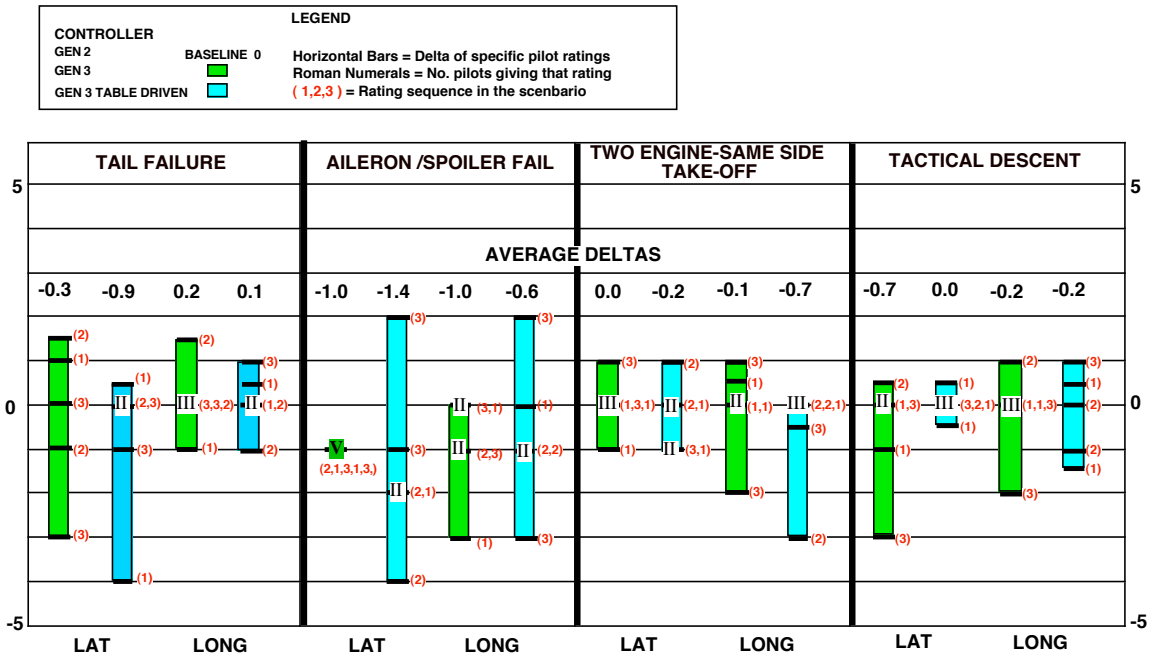

Fig. 4. **Change in Cooper Harper Ratings for Level 2 IFC systems relative to Level 1 IFC system.**

Fig. 5 shows the landing footprint for the entire set of landing scenarios flown. It is important to note that the pilots were not given a specific landing task beyond making a safe landing, but none-the-less, the landing footprint is instructive as to the controllability of the aircraft. The pilots tended to land the aircraft on the left side of the runway because of the crosswind from the left. The Level 2 IFC with the LP allocation routine (green marks) showed a distinctly tighter footprint than the other two controllers. It also had a much smaller standard deviation in the landing sink rate.

Touchdown Circular Error Parameter

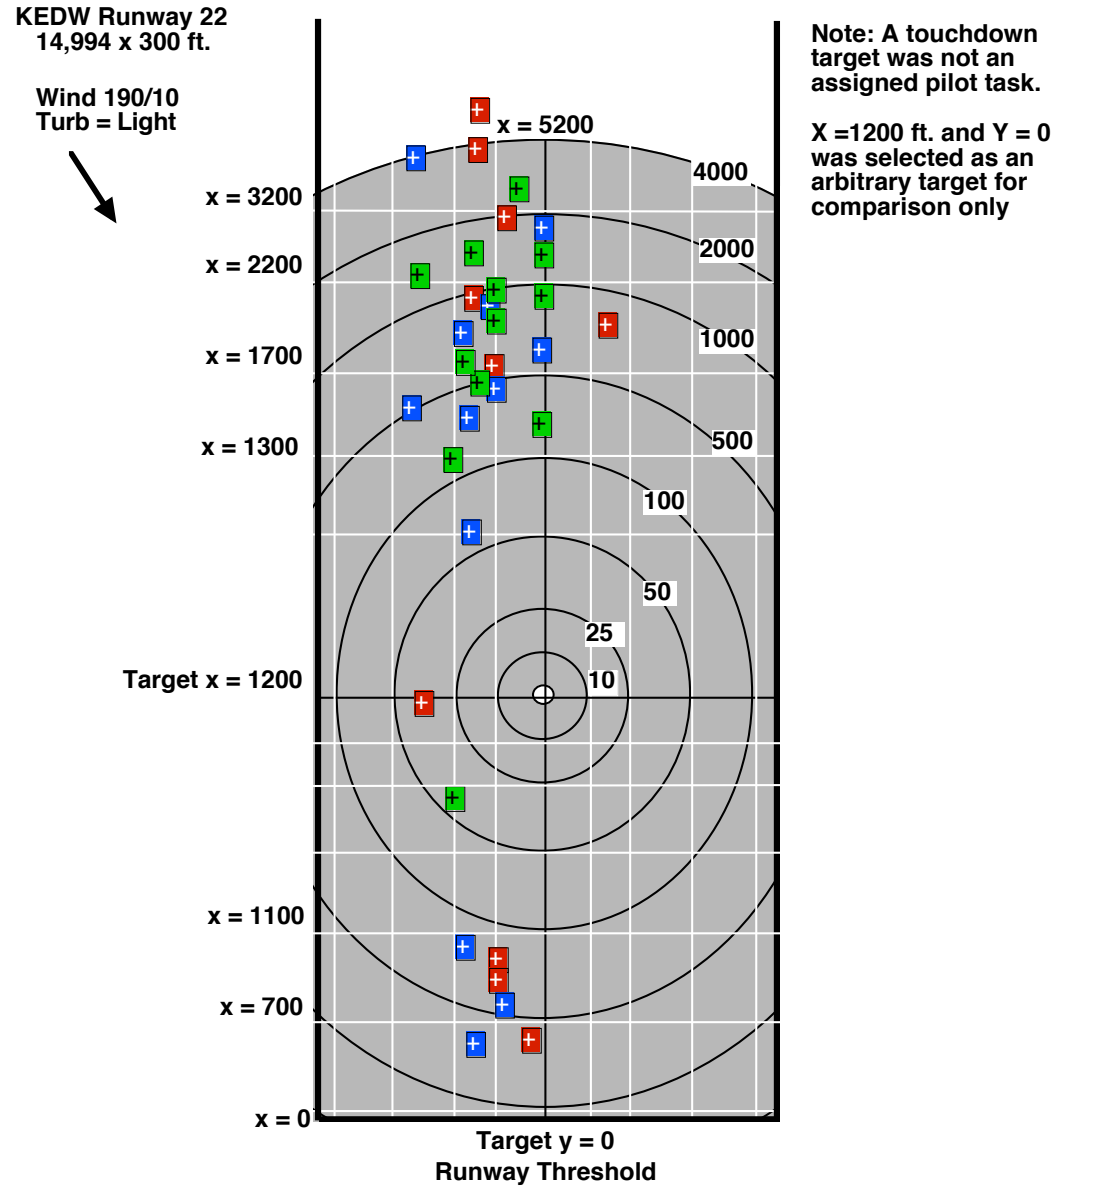

	Touchdown Average			Standard Deviation		
CONTROLLER	X (ft.)	Y (ft.)	Z (Sink Rate) (ft. / sec.)	X	Y	Z
GEN 2	2318.2	-11.6	-3.2	1653.91	12.53	3.20
GEN 3	2079.7	-14.08	-6.3	731.03	10.93	1.74
GEN 3 TABLE DRIVEN	1750.6	-15.1	-7.4	957.76	8.76	2.98

Fig. 5 Touchdown footprint and sink rate upon landing.

Figs. 6 and 7 present sample time history plots from the piloted experiment for the tail failure and tactical descent using the ORC. The data in these figures includes 1: Normalized stick input to the system. Horizontal axis is elapsed time in minutes, 2: Reference model gain parameter (k) for the pitch control of the system, 3: Reference model pitch frequency, 4: Standard approach plot of glide slope deviation, Horizontal axis is Nautical Miles from threshold, 5: Standard approach plot of localizer deviation and 6: Pitch rate error in degrees per second. It is interesting to see the adaptation in the reference model parameters and the overall performance of the piloted aircraft for the landing scenarios. In particular, the coupled failure scenario shows that the ORC approach adapts the parameters only when necessary to accommodate larger errors. This is seen in subplot 2 ands 3 of Fig. 7 in which the adaptation begins only after the third failure signaling that the first three failures did not need any changes to the reference model.

At approximately 160 seconds into the data run (see Figs. 8 and 9), all ailerons and spoilers are failed at their idle position for the wing failure case (Fig. 8) and both rudders and all elevators and stabilizer are failed at their idle position for the tail failure case (Fig. 9). The flight control software must achieve its roll and pitch control, for the respective failures, using nontraditional deployment of the remaining surfaces. The adaptive critic correctly notices the degradation in performance, and rapidly changes gain and frequency values for the reference model. Subsequent control inputs allow the critic to refine these values.

It should be noted that the failure scenarios chosen in this experiment were meant to be dramatic tests of the response of the control system on each axis. As such, the aircraft performance was generally degraded to the point that the adaptive critic typically adjusted the reference model gains and frequencies to the lower limit. It would be expected that for more typical in-flight failures such as single jammed surfaces that the adjustments to the reference model would not be as extreme.

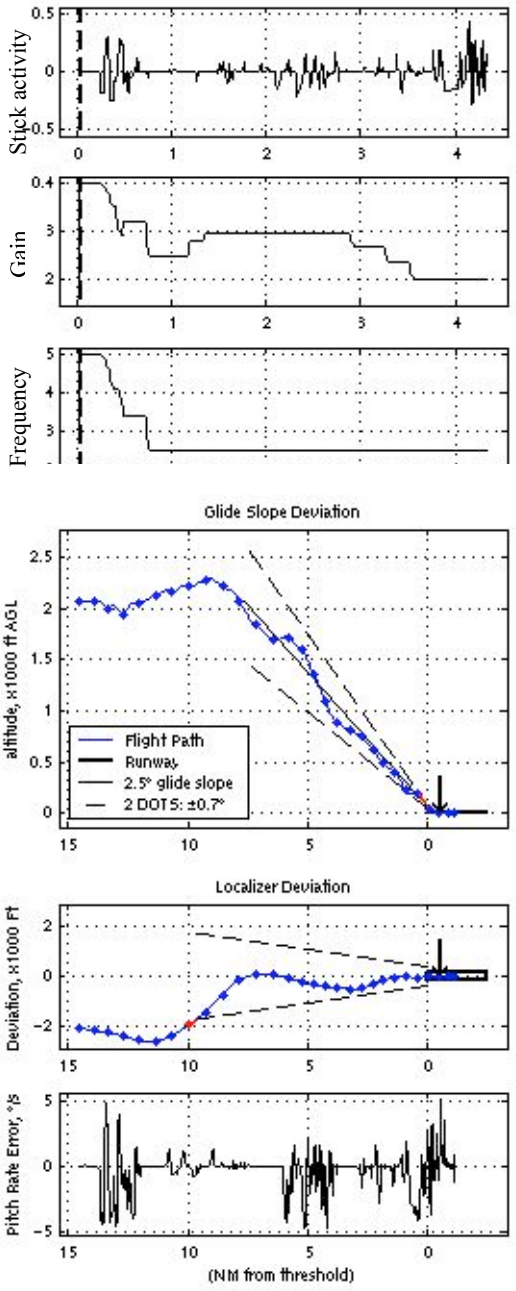

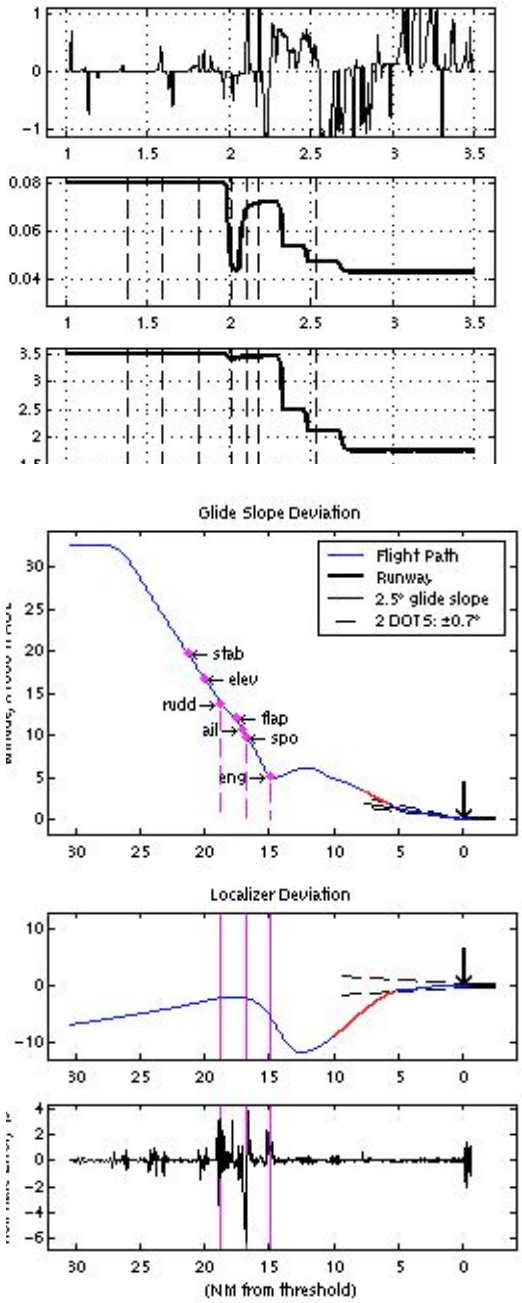

Fig. 6. Data taken during piloted simulation testing (dataset 20030512/122330) using ORC configuration of the critic; the pilot flew a landing with a full tail hydraulic failure (Scenario #1, Table 1). Subplot 1: Normalized stick input to the system. Horizontal axis is elapsed time in minutes. Subplot 2: Reference model gain parameter (k) for the pitch control of the system; Subplot 3: Reference model pitch frequency. Subplot 4: Standard approach plot of glide slope deviation. Horizontal axis is Nautical Miles from threshold. Subplot 5: Standard approach plot of localizer deviation. Subplot 6: Pitch rate error in degrees per second.

Fig. 7. Data taken during piloted simulation testing (dataset 20030512/133532) using ORC configuration of the critic; the pilot flew an assault descent through to a landing with a series of cascaded failures (Scenario #4, Table 1). Subplot 1: Normalized stick input to the system. Horizontal axis is elapsed time in minutes. Subplot 2: Reference model gain parameter (k) for the pitch control of the system; Subplot 3: Reference model pitch frequency. Subplot 4: Standard approach plot of glide slope deviation. Horizontal axis is Nautical Miles from threshold. Subplot 5: Standard approach plot of localizer deviation. Subplot 6: Pitch rate error in degrees per second.

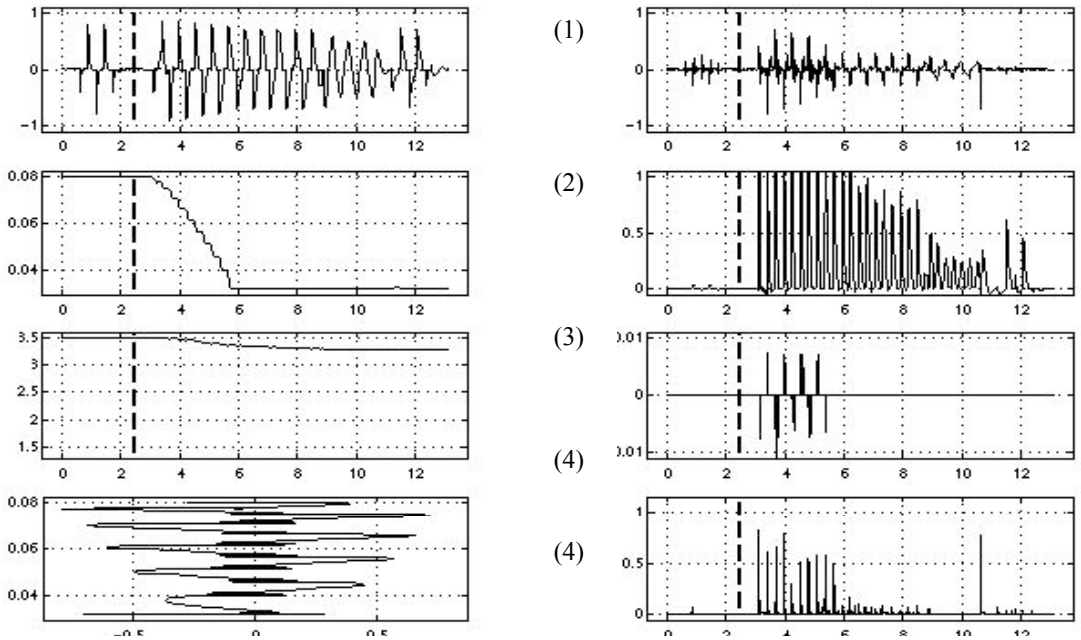

Fig. 8. A synthetic exercise with many doublets in the roll axis; full wing hydraulic failure induced at the vertical dashed mark. Subplot 1 (Left): Normalized stick input to the system. Horizontal axis is elapsed time in minutes. Subplot 2 (Left): Reference model gain parameter (k) for the roll control of the system; Subplot 3 (Left): Reference model roll frequency. Subplot 4 (Left): vertical axis represents the reference model gain (k) and the horizontal axis presents the error in the roll rate. Subplot 1 (Right): Normalized rate error in roll axis. Horizontal axis is elapsed time in minutes. Subplot 2 (Right): Partial derivative of J with respect to the gain, k. Subplot 3 (Right): Excess control demanded in roll axis. Subplot 4 (Right): Square of the error in the adaptive critic network.

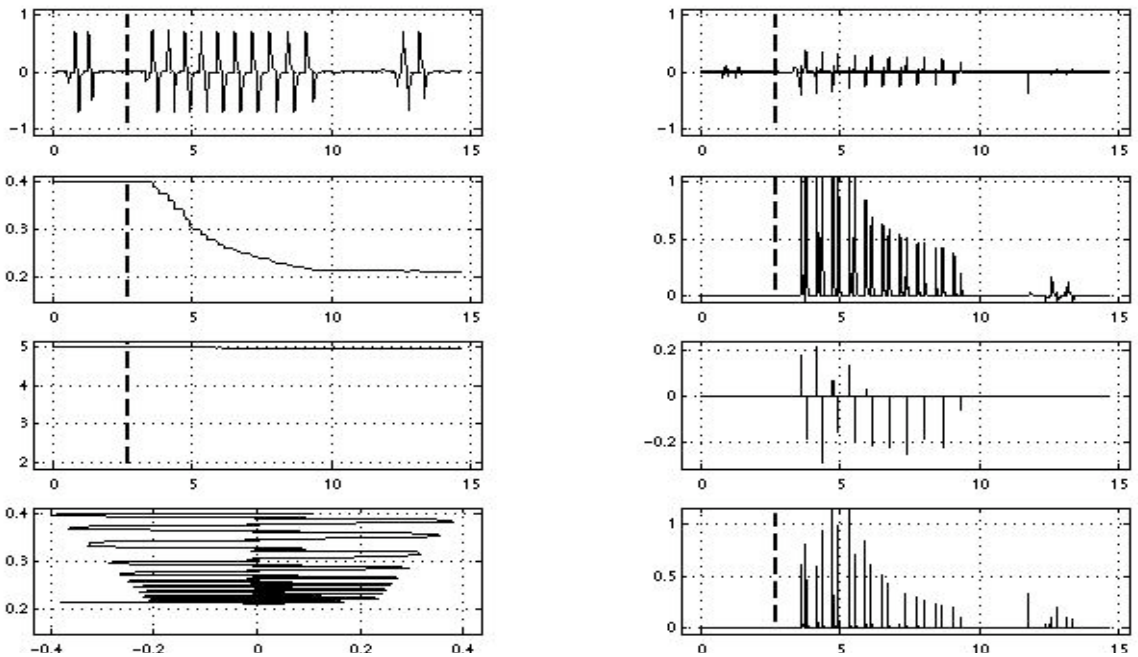

Fig. 9. Pitch axis plots for a full tail hydraulic failure (Scenario 1). Descriptions are the same as for Fig. 8.

VII. Issues for Flight Test Verification

The Level 1 IFC (without control reallocation) is currently scheduled to be flight tested in a highly-modified McDonnell-Douglas NF-15B Eagle at NASA-Dryden in the spring of 2005. This aircraft has been modified from a standard F-15 configuration to include canard control surfaces, thrust vectoring nozzles, and a digital fly-by-wire flight control system. Failures will be simulated by locking out a stabilator control surface as well as through offsetting the canard surfaces relative to their normal schedules. Several issues have been key concerns in the development of this flight test. In particular, transient free reversion to the standard F-15B flight control system in case of unanticipated problems, structural loads models and limitations of the aircraft, reduced stability margins when aeroservoelastic filters were incorporated and validation and verification of the online learning neural networks are all of concern to the flight test effort.

VIII. Conclusions

This paper presented a Level 2 intelligent control architecture that utilizes an adaptive critic to tune the parameters of a reference model, which is then used to define the angular rate commands for a Level 1 intelligent controller. The goal of this on-going research is to improve the performance of the C-17 under degraded conditions such as control failures and battle damage. A simplified version of the adaptive critic was used for the piloted study.

The results presented in the previous sections demonstrate the effectiveness of the neural adaptive flight control system controlling a transport-class vehicle under a wide range of failure conditions. The generic system can also help to reduce the high cost associated with avionics development since it does not require gain-scheduling or explicit system identification. In general, the results demonstrate that under normal flight conditions, the neural system can achieve performance, which is comparable to the aircraft's conventional system. The neural flight control system can also provide additional control authority under damaged or failure conditions. Pilot ratings using a motion based simulation facility showed good handling quality improvement over a Level 1 IFC system.

The choice of control reallocation technique can significantly improve damage adaptation under various failure conditions. Minimization of dead-bands is key to producing good handling qualities, and should be carefully considered in future research. The pilots generally preferred the table-driven scheme for the simple axis-by-axis failure scenarios with control dead-bands. In the final, highly coupled failure, the pilots generally preferred the LP allocation scheme. This suggests that integration of parameter identification techniques, vehicle health monitoring information or incorporation of control surface blending into the cost function will distinctly improve the handling qualities of the LP allocation scheme, and should be pursued in future research.

Adaptive critic approaches are ideally suited for outer-loop trajectory management. Challenges going forward will be to generalize this technology for higher order reference models. Stability of the first order model was easy to ascertain by limiting the parameters. For higher order models, efforts are underway in deriving adaptation laws that guarantee stability.

One of the issues that came up during manned flight was the concept of "control harmony". When reference model gains are modified in one axis and not proportionately in the other, pilots experience a bigger coupling effect at the stick input. Efforts are underway to quantify this effect and to derive certain adaptation heuristics to avoid this man-machine interaction issue. Another issue was the large command deflections (such as split elevator) of non-traditional surfaces required to control the aircraft. The results also imply that aircraft structural design must be carefully evaluated when utilizing control surfaces in non-traditional manners and that structural limitations should be incorporated into the control reallocation.

References

[1]National Transportation Safety Board, Aircraft Accident Report, PB90-910406, NTSB/ARR-90/06, United Airlines Flight 232, McDonnell Douglas DC-10, Sioux Gateway Airport, Sioux City, Iowa, July 1989.

[2]Burcham, F. W., Jr., Burken, John, "Flight testing a propulsion-controlled aircraft emergency flight control system on an F-15 airplane", AIAA-1994-2123

[3]Burken, John J., Burcham, Frank W., "Flight-Test Results of Propulsion-Only Emergency Control System on MD-11 Airplane" J. Guidance, Controls and Dynamics, 1997, 0731-5090 Vol. 20 No. 5 (980-987).

[4]Rysdyk, Rolf T., and Anthony J. Calise, "Fault Tolerant Flight Control via Adaptive Neural Network Augmentation", AIAA 98-4483, August 1998

[5]Kim, B., and Calise, A., "Nonlinear Flight Control Using Neural Networks", *AIAA Journal of Guidance, Navigation, and Control*, Vol. 20, No. 1, 1997.

[6]Kaneshige, J. and Gundy-Burlet, K., "Integrated Neural Flight and Propulsion Control System", AIAA-2001-4386, August 2001.

[7]Gundy-Burlet, K, Krishnakumar, K., Limes, G, & Bryant, D., Control Reallocation Strategies for Damage Adaptation in Transport Class Aircraft, AIAA 2003-5642, August, 2003.

[8]KrishnaKumar, K., Limes, G., Gundy-Burlet, K., Bryant, D., "An Adaptive Critic Approach to Reference Model Adaptation". AIAA 2003-5790, August 2003.

[9]Krishnakumar, K., Levels of Intelligent Control: A Tutorial, New Orleans, LA, 1997.

[10]Neidhoefer, J. C.; Krishnakumar, K.; Intelligent control for near-autonomous aircraft missions, *IEEE Transactions on Systems, Man and Cybernetics*, Part A, Volume: 31 Issue: 1 , 2001 pp: 14 -29.

[11]Werbos, P. J., "Neuro Control and Supervised Learning: An Overview and Evaluation", *Handbook of Intelligent Control*, White D. A., Sofge, D. A., (eds), Van Nostrand Reinhold, NY, 1992.

[12]Prokhorov, D., Wunsch, D. C., "Adaptive Critic Designs", IEEE transactions on Neural Networks, Vol. 8, No. 5, September, 1997.

[13]Kulkarni, N.V.; KrishnaKumar, K., "Intelligent Engine Control Using an Adaptive Critic", pp: 164- 173, *IEEE Transactions on Control Systems Technology*, Volume: 11, Issue: 2, 2003.

[14]Balakrishnan, S. N., Biega, V., "Adaptive Critic Based Neural Networks for Aircraft Optimal Control", *Journal of Guidance, Control and Dynamics*, Vol. 19, No. 4, July-August, 1996.

[15]Smith, G. A., Meyer, G. "Aircraft automatic digital flight control system with inversion of the model in the feed-forward path" AIAA-1984-2627

[16]Norgaard, M., Jorgensen, C., and Ross, J., "Neural Network Prediction of New Aircraft Design Coefficients", NASA TM-112197, May 1997.

[17]Totah, J. J., Kinney, D. J. Kaneshige, J. T. and Agabon, S. "An Integrated Vehicle Modeling Environment", AIAA 99-4106, August 1999.

[18]Blake, M. W., "The NASA Advanced Concepts Flight Simulator: A Unique Transport Aircraft Research Environment", AIAA-96-3518-CP.

[19]Harper, R. P., Jr., Cooper, G. E. "Handling Qualities and Pilot Evaluation" *J. Guidance, Control, and Dynamics* 1986 0731-5090 vol.9 no.5 (515-529)